

EV251222055

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Method And Apparatus For Authenticating An Open  
System Application To A Portable IC Device**

Inventor(s):

John DeTreville

ATTORNEY'S DOCKET NO. MS1-282USC12

1      **RELATED APPLICATIONS**

2      This is a continuation of Application No. 09/287,699, filed April 6, 1999,  
3      entitled "Method and Apparatus for Authenticating an Open System Application to  
4      a Portable IC Device", which is hereby incorporated by reference, and which is a  
5      continuation-in-part of Application No. 09/266,207, filed March 10, 1999, entitled  
6      "System and Method for Authenticating an Operating System to a Central  
7      Processing Unit, Providing the CPU/OS with Secure Storage, and Authenticating  
8      the CPU/OS to a Third Party", which claims the benefit of U.S. Provisional  
9      Application No. 60/105,891, filed October 26, 1998, entitled "System and Method  
10     for Authenticating an Operating System to a Central Processing Unit, Providing  
11     the CPU/OS with Secure Storage, and Authenticating the CPU/OS to a Third  
12     Party".

13      **TECHNICAL FIELD**

14      This invention relates to computer-implemented authentication systems.  
15      More particularly, the invention relates to authentication of an application, running  
16      on an open system, to a portable IC device.

17      **BACKGROUND OF THE INVENTION**

18      Computers are finding more and more uses in a wide variety of fields and  
19      locations. The ability to obtain ever-increasing performance at an ever-decreasing  
20      price is expanding the fields where computers can be used. These reduced costs  
21      make "public" computers more and more plausible. Public computers refer to  
22      computers that are set up and generally available for use by the public, such as in a  
23      hotel room, in a kiosk at an airport or shopping mall, in a store (e.g., a department  
24      store), or in a public library.

1 or grocery store), etc. Such public computers may be interconnected with other  
2 computers, such as other public computers via a local area network (LAN) or other  
3 public and/or non-public computers via a wide area network (WAN) such as the  
4 Internet, or alternatively may be stand-alone computers.

5 The public accessibility to such computers, as well as their potential  
6 interconnectivity, makes each computer an “open system”. An open system refers  
7 to a computer that is accessible to multiple individuals and/or other computers,  
8 some of which cannot be trusted with users’ private information. Open systems  
9 are vulnerable to a wide variety of attacks intended to compromise the integrity of  
10 the systems and reveal users’ private information. Such attacks can come from  
11 other computers via a network (e.g., the Internet), or alternatively from other users  
12 of the systems.

13 Public computers are particularly appealing for use with portable integrated  
14 circuit (IC) devices such as smart cards. A smart card is a small card, roughly the  
15 size of a typical credit card, that includes a microprocessor, memory, and an  
16 input/output (I/O) interface. Smart cards can be programmed to maintain any type  
17 of information. Examples of such information include private financial  
18 information (such as checking or savings account number, credit card numbers,  
19 and personal identification numbers (PINs)), as well as private identification  
20 information (such as a social security number or digital signature).

21 Unfortunately, when public computers are vulnerable to attack, so too are  
22 the smart cards that interface with the computers. A public computer could be  
23 executing, unbeknownst to the user, a “rogue” application that accesses private  
24 information on the smart card and subsequently takes various unauthorized  
25 actions. Examples of such unauthorized actions include charging goods or

1 services to a particular account number and signing the smart card owner's  
2 signature to the charges, transferring money out of checking or savings accounts,  
3 etc. Another type of rogue application executing on the public computer could be  
4 an "impostor" of a legitimate program. For example, a public computer may  
5 include a banking program that allows users, upon providing the appropriate  
6 account numbers from their smart card, to access their current account status or  
7 purchase goods and services. A rogue application may pretend to be the banking  
8 application in order to receive the account numbers provided by the smart card, at  
9 which point various unauthorized actions could be taken by the rogue application.

10 Similarly, a rogue OS (operating system) might intercept a PIN (Personal  
11 Identity Number) or other smart card password entered on the open system's  
12 keyboard, or might intercept communications between the smart card and the  
13 application operating under the OS's control on the open system.

14 One solution that protects private information from rogue applications is to  
15 include, as part of the portable IC device, a display in order to display the requests  
16 being signed by the smart card on behalf of the user, a keyboard in order to allow  
17 the user to enter PINs and to accept or reject requests, and its own clock and  
18 battery supply to provide defense against various other attempts to obtain the  
19 private information. However, this solution provides a rather bulky and expensive  
20 "portable" IC device that is too costly to produce on a mass scale.

21 This invention addresses these disadvantages, providing an improved way  
22 to maintain the security of private information on a portable IC device.

1      **SUMMARY OF THE INVENTION**

2      The invention provides for authentication between an open system and a  
3      portable IC device that can be coupled to the open system. Private or otherwise  
4      sensitive or protected information that is maintained on the portable IC device is  
5      unlocked and made available only to an application, executing on the open system,  
6      that can prove to the portable IC device that it is trustworthy. The trustworthy  
7      application will maintain the security of the private information and will not  
8      misuse the information.

9      According to one aspect of the invention, a secure communication channel  
10     between the open system and the portable IC device is established. An application  
11     desiring access to the information on the portable IC device then authenticates  
12     itself to the portable IC device, proving that it is trustworthy. Once such  
13     trustworthiness is proven, the portable IC device authenticates itself to the  
14     application. Once such two-way authentication has been completed, trusted  
15     communication between the open system and the portable IC device can proceed.

16     According to one aspect of the invention, the open system uses an  
17     “authenticated boot” methodology to authenticate applications executing on the  
18     system. In the authenticated boot methodology, certificates of authenticity can be  
19     provided by the operating system, the processor, and the computer. The operating  
20     system can further provide certificates authenticating particular applications  
21     executing on the open system. A chain of such certificates can then be provided to  
22     the portable IC device, proving the authenticity of the applications.

23     According to another aspect of the invention, the open system uses a  
24     “curtaining” or “curtained code” methodology to authenticate applications  
25     executing on the system. In the curtaining methodology, an application can be

1       executed in a secure manner by the open system, ensuring that no other  
2       applications can access the data being used by the secure application unless  
3       explicitly authorized. A security manager, responsible for handling secure sections  
4       of memory, can provide a certificate that a particular application is executing in a  
5       secure section of memory, thereby proving the authenticity of the application.

6

7       **BRIEF DESCRIPTION OF THE DRAWINGS**

8       The present invention is illustrated by way of example and not limitation in  
9       the figures of the accompanying drawings. The same numbers are used  
10      throughout the figures to reference like components and/or features.

11      Fig. 1 is a block diagram illustrating an exemplary system incorporating the  
12      invention.

13      Fig. 2 is a block diagram illustrating an exemplary computer and portable  
14      IC device in accordance with the invention.

15      Fig. 3 shows general components in an exemplary open system computer.

16      Fig. 4 shows an example of a signed boot block.

17      Fig. 5 shows steps in a method for performing an authenticated boot  
18      operation.

19      Fig. 6 shows an exemplary structure of a boot log.

20      Fig. 7 is an example of a symbolic map of a memory space using  
21      curtaining.

22      Fig. 8 is a block diagram showing relevant parts of a processor that can be  
23      used in an open system computer in accordance with the invention.

24      Fig. 9 is a flowchart illustrating an exemplary process of providing  
25      curtained execution protection in a processor.

1       Fig. 10 details an example of hardware added to a computer for defining  
2       restricted-access storage as a set of secure memory pages or regions within the  
3       same address space as that of system memory.

4       Fig. 11 diagrams components employed in an exemplary open system  
5       computer.

6       Fig. 12 is a flowchart illustrating an exemplary process for activating a  
7       secure loader in accordance with the invention.

8       Fig. 13 is a flowchart illustrating an exemplary process for two-way  
9       authentication in accordance with the invention.

10      Fig. 14 is a flowchart illustrating an exemplary process for a computer  
11      application to authenticate itself to a portable IC device in accordance with the  
12      invention.

13      Fig. 15 is a flowchart illustrating an exemplary process for a portable IC  
14      device to authenticate itself to a computer system in accordance with the  
15      invention.

16

## 17      **DETAILED DESCRIPTION**

18      The discussion herein assumes that the reader is familiar with cryptography.  
19      For a basic introduction of cryptography, the reader is directed to a text written by  
20      Bruce Schneier and entitled “Applied Cryptography: Protocols, Algorithms, and  
21      Source Code in C,” published by John Wiley & Sons with copyright 1994 (or  
22      second edition with copyright 1996).

1      **General System Structure**

2      Fig. 1 is a block diagram illustrating an exemplary system incorporating the  
3      invention. Generally, the system includes multiple open system computers 102  
4      and 104 coupled to a data communications network 106. The communications  
5      network 106 comprises a public and/or private network(s), such as the Internet, a  
6      local area network (LAN) or wide area network (WAN). Additional open system  
7      computers 108 or server computers 110 may also be coupled to communications  
8      network 106, enabling communication among computers 102, 104, 108, and 110.  
9      Each of the open system computers 102-108 can be a public computer that is  
10     generally accessible by the public.

11     Computers 102 and 104 include access ports 112 and 114, respectively.  
12     Access ports 112 and 114 allow a portable integrated circuit (IC) device, such as  
13     device 116, to be communicably coupled to computers 102 and 104 (e.g., device  
14     116 may be inserted into ports 112 and 114). This coupling can be accomplished  
15     in any of a variety of conventional manners. Examples of such couplings include  
16     physical couplings (e.g., electrical contacts on the device 116 that are aligned,  
17     when the card is inserted into ports 112 or 114, with electrical contacts in ports  
18     112 and 114), wireless couplings (e.g., an infrared (IR) connection), etc.

19     By coupling an IC device 116 to one of the computers 102 or 104, a user is  
20     able to access his or her private information on device 116, or private information  
21     maintained on another computer or server, or perform other restricted actions. For  
22     example, a user may access his or her bank account, or alternatively order  
23     products or services and “sign” the order using IC device 116. Such uses for  
24     portable IC devices are well-known to those skilled in the art and thus will not be  
25     discussed further except as they pertain to the invention.

1      Portable IC device 116 can be any of a wide variety of portable devices.  
2      One such portable device is a smart card, having approximately the same  
3      dimensions as a conventional credit card. Device 116 can alternatively be  
4      implemented as other types of portable devices, such as jewelry (e.g., a pendant,  
5      ring, or necklace), electronic devices (e.g., a pocket organizer or cellular phone),  
6      watches, etc.

7      When IC device 116 is coupled to one of computers 102 or 104, the IC  
8      device 116 authenticates itself to the computer in order to ensure the computer that  
9      device 116 itself is indeed present. This authentication of IC device 116 also  
10     ensures to the computer that the proper user of the device 116 is present (e.g., the  
11     device 116 has not been stolen). Additionally, in accordance with the invention  
12     one or more applications running on the computer are authenticated to IC device  
13     116. The authentication of the application(s) ensures to IC device 116 that the  
14     application is trusted to not improperly access or use confidential information  
15     obtained from IC device 116.

16     One or more software application(s) running on computer 102 or 104 can  
17     be authenticated to portable IC device 116 using different methodologies. One  
18     way to support such authentication is referred to as “authenticated boot”. Using  
19     the authenticated boot methodology, the operating system on the computer is able  
20     to prove its identity to the microprocessor and thereby certify that it is trusted.  
21     The operating system can then certify other applications that are to be trusted.  
22     Another way to support such authentication is referred to as “curtaining” or  
23     “curtained code”. Using the curtaining methodology, trusted applications can be  
24     executed in a secure manner regardless of the trustworthiness of the operating  
25     system. A security manager coordinates such execution, and can provide

1 certificates proving that particular applications are executing in a secure manner.  
2 Both the authenticated boot and curtaining methodologies are discussed in more  
3 detail below.

4 Fig. 2 is a block diagram illustrating an exemplary computer and portable  
5 IC device in accordance with the invention. The computer 118 represents an open  
6 system computer, such as computer 102 or 104 of Fig. 1. Computer 118 includes  
7 a processor 120 and memory 122 storing an operating system (OS) 123 and  
8 multiple applications 124, coupled together by an internal or external bus 125.  
9 Memory 122 represents any of a variety of conventional volatile and/or  
10 nonvolatile memory or storage components, such as read only memory (ROM),  
11 random access memory (RAM), magnetic storage (e.g., tapes, or floppy or hard  
12 disks), optical disks (e.g., CD-ROMs or DVDs), etc. Any one or more of  
13 applications 124 or operating system 123 may request that portable IC device 116  
14 unlock itself, thereby allowing the requesting application or operating system to  
15 access private information stored on the portable IC device.

16 Portable IC device 116 includes a processor 126 and memory 128 having a  
17 data storage section and an authentication application 130, coupled together by an  
18 internal bus 131. Memory 128 represents any of a variety of conventional  
19 nonvolatile storage components, such as ROM or flash memory. Alternatively, if  
20 portable IC device 116 were to have a separate power source (e.g., a small  
21 battery), memory 128 could also include volatile memory. Data storage 129  
22 includes the user's private information. Authentication application 130 is an  
23 application that interacts with the application(s) 124 and/or operating system 123  
24 of computer 118 to determine whether to unlock the portable IC device 116 and  
25 thereby make the information in data storage 129 available to the requesting

1 application of computer 118. Portable IC device 116 will not provide any data in  
2 data storage 129 to applications 124 or operating system 123 until the requesting  
3 application or operating system is authenticated. An optional indicator light 132  
4 or other signaling device is used to inform the user that the computer 116 has  
5 appropriately authenticated itself to the portable IC device 116. The operation of  
6 authentication application 130 is discussed in more detail below.

7

8 **Authenticated Boot**

9 Fig. 3 shows general components in an exemplary computer 118 of Fig. 2.  
10 They include a central processing unit (CPU) 134, nonvolatile memory 136 (e.g.,  
11 ROM, disk drive, CD ROM, etc.), volatile memory 138 (e.g., RAM), a network  
12 interface 140 (e.g., modem, network port, wireless transceiver, etc.), and an OEM  
13 certificate 141 whose use is described below. The computer 118 may also include  
14 a sound system 142 and/or a display 144. These components are interconnected  
15 via conventional busing architectures, including parallel and serial schemes (not  
16 shown).

17 The CPU 134 has a processor 146 and may have a cryptographic  
18 accelerator 148. The CPU 134 is capable of performing cryptographic functions,  
19 such as signing, encrypting, decrypting, and authenticating, with or without the  
20 accelerator 148 assisting in intensive mathematical computations commonly  
21 involved in cryptographic functions.

22 The CPU manufacturer equips the CPU 134 with a pair of public and  
23 private keys 150 that is unique to the CPU. For discussion purposes, the CPU's  
24 public key is referred to as " $K_{CPU}$ " and the corresponding private key is referred to  
25 as " $K_{CPU}^{-1}$ ". Other physical implementations may include storing the key on an

1 external device to which the main CPU has privileged access (where the stored  
2 secrets are inaccessible to arbitrary application or operating system code). The  
3 private key is never revealed and is used only for the specific purpose of signing  
4 stylized statements, such as when responding to challenges from a portable IC  
5 device, as is discussed below in more detail.

6 The manufacturer also issues a signed certificate 152 testifying that it  
7 produced the CPU according to a known specification. Generally, the certificate  
8 testifies that the manufacturer created the key pair 150, placed the key pair onto  
9 the CPU 134, and then destroyed its own knowledge of the private key " $K_{CPU}^{-1}$ ".  
10 In this way, nobody but the CPU knows the CPU private key  $K_{CPU}^{-1}$ ; the same key  
11 is not issued to other CPUs. The certificate can in principle be stored on a  
12 separate physical device but still logically belongs to the processor with the  
13 corresponding key.

14 The manufacturer has a pair of public and private signing keys,  $K_{MFR}$  and  
15  $K_{MFR}^{-1}$ . The private key  $K_{MFR}^{-1}$  is known only to the manufacturer, while the  
16 public key  $K_{MFR}$  is made available to the public. The manufacturer certificate 152  
17 contains the manufacturer's public key  $K_{MFR}$ , the CPU's public key  $K_{CPU}$ , and the  
18 above testimony. The manufacturer signs the certificate using its private signing  
19 key,  $K_{MFR}^{-1}$ , as follows:

20  
21 Mfr. Certificate =  $(K_{MFR}, \text{Certifies-for-Boot}, K_{CPU})$ , signed by  $K_{MFR}^{-1}$

22  
23 The predicate "certifies-for-boot" is a pledge by the manufacturer that it  
24 created the CPU and the CPU key pair according to a known specification. The  
25 pledge further states that the CPU can correctly perform authenticated boot

1 procedures, as are described below in more detail. The manufacturer certificate  
2 152 is publicly accessible, yet it cannot be forged without knowledge of the  
3 manufacturer's private key  $K_{MFR}^{-1}$ .

4 Another implementation in which a 'chain of certificates' leading back to a  
5 root certificate held by the processor manufacturer is also acceptable.

6 Similarly, the OEM (Original Equipment Manufacturer), the manufacturer  
7 of the computer as distinguished from the manufacturer of the processor, may  
8 provide an OEM certificate 141 that certifies that the design of the computer  
9 external to the processor does not include various known attacks against the secure  
10 operation of the processor. The OEM also has a pair of public and private signing  
11 keys,  $K_{OEM}$  and  $K_{OEM}^{-1}$ . The OEM certificate is signed using the private key  $K_{OEM}^{-1}$   
12 <sup>1</sup> analogous to the manufacturer's certificate 152 being signed by the processor  
13 manufacturer.

14 The CPU 134 has an internal software identity register (SIR) 154, which is  
15 cleared at the beginning of every boot. The CPU executes an opcode  
16 "BeginAuthenticatedBoot" or "BAB" to set an identity of a corresponding piece of  
17 software, such as operating system 160, and stores this identity in the SIR; the  
18 boot block of the operating system (described below) is atomically executed as  
19 part of the BAB instruction. If execution of the BAB opcode and the boot block  
20 fails (e.g., if the execution was not atomic), the SIR 154 is set to a predetermined  
21 false value (e.g., zero). This process is described below in more detail with  
22 reference to Fig. 5.

23 The CPU 134 also utilizes a second internal register (LOGR) 156, which  
24 holds contents produced as a result of running a LOG operation. This operation,  
25 as well as the register, is described below in more detail.

1        The CPU 134 also maintains a “boot log” 158 to track software modules  
2 and programs that are loaded. In one implementation, the boot log 158 is a log in  
3 an append-only memory of the CPU that is cleared at the beginning of every boot.  
4 Since it consumes only about a few hundred bytes, the boot log 158 can be  
5 comfortably included in the main CPU. Alternatively, the CPU 134 can store the  
6 boot log 158 in volatile memory 138 in a cryptographic tamper-resistant container.

7        A further implementation is by means of a software module that allows  
8 each section of the booting operating system to write entries into the boot log that  
9 cannot be removed by later components without leaving evidence of tampering.  
10 Yet alternatively, the SIR can hold a cryptographic digest of a data structure  
11 comprising the initial boot block and the subsequent contents of the boot log. The  
12 operation of appending to the boot log (call this operation “Extend”) replaces the  
13 SIR with the hash of the concatenation of the SIR and the entry being appended to  
14 the boot log. A straightforward implementation of this operation may be seen to  
15 modify the SIR. Note, however, that the operating system, when booting, can  
16 choose to add elements to the boot log without loading the corresponding  
17 components, and so a more privileged combination of software components can  
18 impersonate a less privileged one. This allows the controlled transfer of secrets  
19 across privilege levels. In this approach, software will keep its own plaintext copy  
20 of the boot log entries, along with the initial value of the SIR following boot, and  
21 this plaintext copy is validated by knowledge of the current composite SIR.

22        As an optimization, regardless of the implementation of the boot log, the  
23 OS may choose not to extend the boot log with the identities of certain software  
24 components, if these components are judged to be as trustworthy as the OS itself,  
25

1 or if they will execute only in a protected environment from which they will be  
2 unable to subvert operation.

3 The operating system (OS) 160 is stored in the memory 136 and executed  
4 on the CPU 134. The operating system 160 has a block of code 162 used to  
5 authenticate the operating system on the CPU during the boot operation. The boot  
6 block 162 uniquely determines the operating system, or class of operating systems  
7 (e.g. those signed by the same manufacturer). The boot block 162 can also be  
8 signed by the OS manufacturer.

9 Fig. 4 shows an example of a signed boot block 180 created by signing the  
10 block of code 162. It contains the BeginAuthenticatedBoot opcode 182, a length  
11 184 specifying the number of byte in the block of code, the code 162, a signature  
12 186, and a public key 188 used to verify the signature 186. The boot block will  
13 also contain as a constant or set of constants, keys, or other information 190 that is  
14 used to validate the subsequent operating system components (for instance a  
15 public key or keys). In this implementation, the CPU will set the SIR to the public  
16 key of the boot block, but only if the boot block code signature is correct for the  
17 stated boot block public key.

18 In an alternative implementation, the SIR is set to the cryptographic hash or  
19 digest of the code and constants that make up the boot block. The signature 186  
20 and public key 188 are then not needed.

21 A key observation of both of these implementations is that no one can boot  
22 an untrusted operating system in which the SIR is set to the value of a trusted  
23 operating system.

24 Once booted the operating system 160 and other selected applications (e.g.,  
25 banking applications) named in an access control list (ACL) by the owner of the

1 computer can set aside space 164 in memory or disk 136 to hold private or  
2 confidential data in a secure manner, without fear of other operating systems or  
3 rogue applications reading the data in the space. The private data is protected by  
4 encryption using a key that is generated based in part upon a seed supplied by an  
5 authenticated and trusted OS, in part by a secret key stored in the CPU, and in part  
6 by the software identity register (SIR). The private data is stored with an ACL  
7 naming the applications that can use the data and the terms under which they can  
8 use it.

9 Software programs 166 (the applications) are also shown stored in memory  
10 136. These programs may interact with the portable IC device 116. Each program  
11 166 has an associated key or digest 168 for unique identification.

12 The authenticated boot process allows any software at any point in the boot  
13 sequence to initiate an authenticated boot.

14 Fig. 5 shows steps in a method for performing an authenticated boot  
15 operation on the operating system 160. These steps are performed by the CPU 134  
16 and OS 160 resident in the computer 118. At step 200, the CPU executes the  
17 BeginAuthenticatedBoot opcode 182 in the signed boot block 180 to set an  
18 identity for the operating system 160. The identity can be a digest of the boot  
19 block's opcodes and data, or the public key 188 corresponding to a signature on  
20 the boot block of the operating system.

21 The BeginAuthenticatedBoot opcode 182 and the boot block 180 execute as  
22 one atomic operation, with the implication that if they execute completely and  
23 correctly, the resulting operating system can be trusted. Measures are taken to  
24 ensure that the CPU is not interrupted and that the boot code that has just been  
25 validated cannot be modified. This can involve locking the memory bus and

1 switching off interrupts. It could also involve having the CPU watch for interrupts  
2 or for writes by other bus agents and invalidate the authenticated boot sequence if  
3 they occur. The BAB opcode 182 can be executed at any time, with one  
4 exemplary time being at the start of the OS loader, right after the OS-selector  
5 executes. An alternative implementation is to provide both a  
6 BeginAuthenticatedBoot (BAB) and an EndAuthenticatedBoot (EAB) instruction.  
7 The BAB instruction computes the secure hash of the boot block and the EAB  
8 instruction sets the SIR if the execution of the boot block was not interrupted or  
9 potentially modified by memory writes from another processor or another bus  
10 master.

11 Execution of the BeginAuthenticatedBoot opcode 182 sets the internal  
12 software identity register 158 to either (1) the OS's identity (i.e., boot block digest  
13 or OS public key 188) if the operation is successful, or (2) zero if some event or  
14 circumstance has potentially subverted operation. Assuming the operation is  
15 successful (i.e., the "yes" branch from step 202), the SIR 154 is now a unique  
16 number or other value that represents the identity of the operating system 160  
17 (step 204). Any two processors running the same operating system will produce  
18 the same SIR. If the BAB opcode operation is unsuccessful (i.e., the "no" branch  
19 from step 202), the SIR is set to zero (step 206).

20 It is noted that different operating systems may be serially booted on the  
21 computer 118. Executing the BAB opcode 182 for different signed OS boot  
22 blocks results in different SIR values. However, it is possible for multiple boot  
23 blocks to result in the same SIR, when desired.

24 At step 208, the CPU 134 fills the first entry on the boot log 158 with the  
25 public key (or digest) of the boot block 162. From now on, any running code can

1 append data to the boot log 158, and it is generally used by code in the boot chain  
2 to identify code versions as they are loaded and executed. As noted earlier,  
3 appending data to the boot log can be simulated by modifying the SIR via the  
4 “Extend” operation.

5 The boot block 162 is free to load the next set of blocks in the boot-chain  
6 (step 210). At step 212, the boot block 162 checks the validity of the modules (by  
7 signature or other means) and loads them so that they can be executed. An  
8 identity for each module is appended to the boot log 158. The OS will also retain  
9 additional information on components that it loads (e.g., version numbers, device  
10 driver IDs, etc.). Loading and executing the code may result in loading more  
11 code, validating it, and executing it, etc. This process continues through to the  
12 loading of device drivers. When the boot sequence is complete, the OS is  
13 operational and the software identity register and the boot log store non-  
14 modifiable data captured during the boot sequence. Loading new device drivers  
15 can be recommenced at any point, possibly causing the operating system to  
16 become less privileged, with the possible termination of access to private data.

17 The CPU can generate a signed certificate containing the boot log data to  
18 attest to the particular operating system (including drivers) that is running. It  
19 could also generate a signed statement containing just the SIR. Fig. 6 shows an  
20 exemplary structure of a boot log 158. It contains a seed field 222 and a block ID  
21 field 224. The block ID field 224 holds identities of the blocks of code that are  
22 loaded and verified on the subscriber unit. The block ID field 224 can hold text or  
23 binary data.

24 The SIR or the seed field 222 holds an authenticated boot key generator  
25 seed. The CPU uses the seed in field 222 to generate keys unique to the OS and

1 processor. Since the first entry of the boot log 158 can only be generated by the  
2 execution of a particular boot block or the holder of the boot block private key, the  
3 keys can only be re-generated by the same OS, or another OS from the same  
4 publisher under control of the publisher.

Once the CPU has derived an appropriate SIR for the operating system, the combination of the CPU and the OS has a unique identification that may be presented to third parties. The computer 118 is thus prepared to communicate with a portable IC device, to specify the CPU and the OS, and prove the identity of the CPU and operating system to the portable IC device.

10 The action of signing a statement of the current value of the SIR can be  
11 generalized into a technique for the operating system to make a signed attestation  
12 of the current value of any arbitrary region of memory and/or a register. In one  
13 implementation, the ATTEST operation is performed as follows:

14

15 ATTEST(Register Name, Region of Memory)

16

17 This operation produces a signed result:

18

19 [K<sub>CPU</sub>, "ATTEST", Register Name, Register Value, Memory Contents]

20 (signed with  $K_{CPU}^{-1}$ )

21

22 The ATTEST operation can be used to sign the current value of the SIR or  
23 any other register (including a log register LOGR, described below). The  
24 processor also signs the data contained in an arbitrary region of memory. This can  
25 be used to include the challenge value or some other signed statement desired by

1 the operating system or application. The ATTEST operation can be used to  
2 provide an implementation of a more general form of OS certificates.

3 The boot log is an append-only record of all or selected components loaded  
4 by the operating system. This can be managed entirely in hardware, but can be  
5 simplified by means of a LOG operation.

6 The LOG operation constructs a secure one way digest of a supplied  
7 parameter, an internal LOGR register 156, and stores the result back in the LOGR  
8 register. At power up, or at processor reset, the LOGR register is set to zero. The  
9 LOGR register can be read, but not written apart from execution of the LOG  
10 operation. The processor can also sign a statement attesting to the current value of  
11 the LOGR register and a supplied challenge. Symbolically:

$$13 \quad \text{LOGR}' = \text{SHA-1}(\text{LOGR}, \text{DATA}) \quad (1)$$

15 where LOGR is the current value of the register, DATA is supplied to the LOGR'  
16 and is the contents of the LOGR register after the LOG operation is performed,  
17 and SHA-1 is an exemplary one way hash function (the Secure Hash Algorithm).

18 The operating system can use the LOG operation to record the digest of  
19 each component as it is loaded. When communicating with a portable IC device,  
20 the ATTEST operation can be used to provide an un-tamperable attestation of all  
21 components loaded into the operating system.

22 In order for the portable IC device to be able to interpret the LOGR value,  
23 the operating system also conveys the digests of all components that made up the  
24 boot log. The portable IC device can then:

- 1     1. Check that all of the components revealed are known and trusted.
- 2     2. Check that the composite value obtained when the digests are combined
- 3         according to equation (1) match that quoted by the microprocessor.
- 4     3. Check that the signature on the quoted statement is valid for the
- 5         microprocessor public key.
- 6     4. Check that the processor certificate and OEM certificate are valid and
- 7         correspond to the processor key used in the quoted statement.

If these conditions are met, the portable IC device can trust the client with the confidential data. If they are not met, then the portable IC device can return a statement of the components that are not trusted so that the administrator of the open system can upgrade the untrusted component(s).

The fundamental requirements of atomicity and privileged access to keys for the microcode that implements authenticated boot can be met in a variety of alternative implementations. In one implementation, components in the chipset may examine the bus to infer operation and permit or deny access to keys depending on the code executing. Components on the chipset can also examine the bus for unauthorized agents writing to protected code, or reading unauthorized secrets.

An agent on the bus can also check for unauthorized interrupts during the execution of the authenticated operations or execution of the boot block.

Similarly, there is no fundamental requirement for the microcode that implements the authenticated boot operations to be physically resident on the microprocessor chip. It could also be stored in ROM, EPROM, or protected flash memory in a physically separate device on the bus.

1        The authenticated boot technique can be implemented by existing CPU  
2 operating modes using code in the computer's BIOS code. The System  
3 Management Mode (SMM), supported by Intel microprocessors, provides for a  
4 region of memory that is inaccessible to normal operating system operation, but  
5 can provide subroutines that operating systems or applications can use. Such  
6 SMM protected memory could be used for the storage of keys and the code that  
7 manages those keys.

8        Hash algorithms and signature schemes can be broken. One example of a  
9 security break is that an attacker finds a second boot block that has the same  
10 identity (same signature, or same digest). If such a boot block is found, then a  
11 different operating system can be booted, and all data security is lost.

12       Greater security can be obtained by combining security schemes. For  
13 instance, the OS-identity can be formed as the concatenation of two or more  
14 digests calculated using different hash algorithms, or the same algorithm applied  
15 to boot block data in a different order (for instance, backwards).

16       In the case of signature based identity, the boot block can be signed several  
17 times using different signature algorithms and keys. Again the software identity  
18 becomes the concatenation of the relevant public keys. This technique also  
19 provides protection against private key compromise. If one of the signing keys is  
20 compromised, not all security is lost.

21       Microprocessor key compromise is inevitable and is traditionally handled  
22 by revocation lists (list of untrusted hosts). However, if the certificates that vouch  
23 for the microprocessor keys never expire, then revocation lists will grow  
24 uncomfortably large. This problem can be ameliorated by finite lifetime processor  
25 certificates and OEM certificates. Portable IC devices will require valid

1 certificates (not expired), and the chip vendors and OEMs (or other trusted parties)  
2 will be required to re-issue certificates for chips and computers still considered in  
3 good standing (not revoked). Note that the certificates are not used when offline,  
4 so machines do not stop working when the certificates expire. However, in online  
5 transactions, an occasional (probably automated) extra step would be to get a new  
6 certificate.

7

### 8 **Curtained Code**

9 Fig. 7 is a symbolic map of a memory space 252 in computer 118 of Fig. 2.  
10 For purposes of illustration, consider it to have a potential size of 4Gbytes, so that  
11 32 bits of address suffice to access all of it. Space 252 can exist in a single  
12 physical memory, or in several different kinds of storage, such as ROM, read/write  
13 RAM, flash RAM, and so forth. Also, partially or totally separate address spaces  
14 are a straightforward extension. Space 252 has three regions or rings 254, 256,  
15 and 258 relevant to the present discussion. Although the information stored in  
16 these rings can be similar to that contained in the rings sometimes used in  
17 processors that employ conventional privilege levels or operational modes, their  
18 mechanism is quite different and independent.

19 Ring 254 is called Ring C or the outer ring, and has only conventional  
20 protection or security against any kind of read or write access by any code located  
21 there or in the other rings in the present system, and normally occupies almost all  
22 of the available address space. All normal user-mode code and data resides in this  
23 ring. The operating system, including the kernel, also resides there. Ring C has  
24 no read or write access to the other two rings.

1        Rings 256 and 258 together comprise the secure or curtained region of  
2 memory. No program code in Ring C has any access to data within them. Ring C  
3 code, can, however, be provided some ability to initiate the execution of code  
4 located there, as described below. Conversely, any code in rings 256 and 258 has  
5 full conventional access to Ring C, including reading and writing data, and  
6 executing program code.

7        Ring 256, also called Ring B, has full access privileges to Ring C, but only  
8 restricted access to ring 258. Thus, it can have both semipermanent storage such  
9 as nonvolatile flash RAM for code routines and volatile read/write memory for  
10 temporary data such as keys. A megabyte or less of the total address range would  
11 likely suffice for Ring B.

12       Ring 258, also called Ring A or the inner ring, in turn has full access to  
13 Rings B and C for both code and data. It can also employ both nonvolatile and  
14 volatile technologies for storing code and data respectively. Its purpose is to store  
15 short loader and verifier programs and keys for authentication and encryption.  
16 The address space required by Ring A is generally much smaller than that of Ring  
17 B. That is, this exemplary embodiment has the Ring A address range within the  
18 address range of Ring B, which in turn lies within the address range of Ring C.  
19 The address ranges of the rings need not be contiguous or lie in a single block. In  
20 order to prevent the access restrictions of the curtained rings from being mapped  
21 away by a processor, the address ranges of Rings A and B can be treated as  
22 physical addresses only. In one embodiment, virtual addresses are conventionally  
23 translated into their corresponding real addresses, and then the restrictions are  
24 interposed at the level of the resulting real addresses. Alternatively, a mechanism  
25 could disable virtual addressing when certain addresses are accessed.

1        In the contemplated area of authentication of rights, it can be desirable to  
2        allow multiple parties to emplace their own separate authentication code and data  
3        that cannot be accessed by any of the other parties. For example, the manufacturer  
4        of the processor, the manufacturer of the computer, the provider of the operating  
5        system, and the provider of trusted application programs may all desire to execute  
6        their own authentication or other security routines and manage their own keys. At  
7        the same time, each party should be able to use code and data in the unsecure Ring  
8        C, and to execute certain routines in the inner Ring A. Dividing Ring B into peer  
9        subrings 260, 262, and 264 permits this type of operation. Region 260, called  
10       Subring B1, has the privileges and restrictions of Ring B, except that it cannot  
11       access subring 262 or 264.

12       Subring B1 can, however, access any part of Ring B that lies outside the  
13       other subrings. In this way, Subring B1 can function as though it were the only  
14       middle ring between Rings A and C. Subrings 262 (B2), and 264 (B3) operate in  
15       the same manner. A typical PC-based system might have three or four subrings,  
16       of 64-128KBytes each. The code in these subrings is normally updated seldom, so  
17       that conventional flash memory is appropriate. Alternatively, the Ring-A loader  
18       could load the code and keys into RAM from an encrypted storage on disk on  
19       demand. Each subring will also require a small amount of scratch RAM, although  
20       rewritable flash memory might be suitable here as well; it might be desirable to  
21       use this for persisting the state of the system after a reboot. For extra flexibility,  
22       the memory available to the curtained memory subsystem can be allocated under  
23       the control of the Ring-A executive code. In order that no untrusted party can  
24       manipulate the memory map to reveal secrets, the map of the subrings in the Ring-  
25

1 B memory is kept in flash storage in curtained memory, under control of the  
2 curtained-memory controller in ring A.

3 In presently contemplated authentication procedures, Ring A code and keys  
4 are loaded under conditions in which protection against snoopers is not necessary;  
5 for example, they can be loaded when the microprocessor is manufactured. This  
6 simple step eliminates any requirement for building any cryptographic capabilities  
7 into the processor itself. Accordingly, Ring A code and keys can be stored in  
8 permanent ROM, with only a few hundred bytes of scratchpad RAM. This Ring A  
9 code is designed to load further curtained code and keys into ring B memory  
10 segments through a physically insecure channel, such as a public network, in such  
11 a manner that an eavesdropper, including even the owner of the target computer,  
12 cannot discover any secret information contained therein. This downloaded code,  
13 operating from the secure memory, then performs the authentication operations  
14 that users require before they will trust their valuable data to the rights-  
15 management software of the system. This bootstrapping procedure permits  
16 building a wide class of secure operations and associated secret keys with greater  
17 security than would be possible in traditional assembly code, even with some form  
18 of authentication routines.

19 However, there are no restrictions on the code that can be loaded into any  
20 of the Ring-B memory areas. Examples of Ring-B code include applications for  
21 key management, applications for interfacing with and accessing information on  
22 portable IC devices, secure storage, signing, and authentication. Further examples  
23 include electronic cash storage, a secure interpreter for executing encrypted code,  
24 modules for providing a software licenses necessary for a piece of software to run.  
25 It is also possible to load only a part of an application, such as a module that

1 communicates with a media player in unsecure memory for reducing software  
2 piracy.

3 The foregoing shows how untrusted code can be prevented from accessing  
4 the contents of a secure memory. The trusted code that is permitted to perform  
5 secure operations and to handle secret data is called curtained code. In other  
6 systems, such code must be executed within a privileged operating mode of the  
7 processor not accessible to non-trusted software, or from a separate secure  
8 processor. In the present invention, however, curtained code can only be executed  
9 from particular locations in memory. If this memory is made secure against  
10 intrusion, then the curtained code can be trusted by third parties. Other features  
11 restrict subversion through attempts at partial or modified execution of the  
12 curtained code.

13 Fig. 8 is a block diagram showing relevant parts of a processor 282 that can  
14 be used in an open system computer, such as computer 118 of Fig. 2. Internal  
15 buses 284 carry data, address, and control signals to the other components of the  
16 processor on the integrated-circuit chip or module. Line 286 carries some of these  
17 signals to and from a bus controller to communicate with an external bus.  
18 Conventional function or execution units 288 perform operations on data from  
19 external memory, from register files 290, from cache 292, from internal  
20 addressable memory 294, or from any other conventional source. Memory 294,  
21 located on the same chip or module as the rest of processor 282, can have a  
22 number of technologies or combinations of technologies, such as dynamic  
23 read/write, read-only, and nonvolatile such as flash. The internal memory in this  
24 implementation partakes of the same address sequence as external system  
25 memory, although it can have or be a part of another sequence. The curtained

1 memory rings can be partly or totally contained in addresses located within  
2 memory 294.

3 Control unit 296 carries out a number of operations for sequencing the flow  
4 of instructions and data throughout the processor; line 298 symbolizes control  
5 signals sent to all of the other components. Interrupt logic 300 receives interrupt  
6 requests and sends system responses via lines 302; in some systems, interrupt  
7 logic is conceptually and/or physically a part of an external bus controller. A  
8 conventional instruction pointer holds the address of the currently executing  
9 instruction. Instruction decoder 304 receives the instruction at this address on line  
10 306, and produces a sequence of control signals 298 for executing various phases  
11 of the instruction. In modern pipelined and superscalar microprocessors, blocks  
12 308 and 304 become very complex as many instructions are in process at the same  
13 time. Their basic functions, however, remain the same for the present purpose.

14 Control unit 296 further includes a specification or map 310 of one or more  
15 address ranges of the memory addresses desired to be curtained. The specification  
16 can be in any desired form, such as logic circuitry, a read-only table of addresses  
17 or extents, or even a small writable or rewritable storage array. If the addresses  
18 are in memories having separate address sequences, additional data specifying the  
19 particular memories can be added to the addresses within each sequence. A  
20 detector or comparator 312 receives the contents of instruction pointer 308 and the  
21 curtained-memory map 310. A curtained memory having multiple rings, subrings,  
22 or other levels can have a separate specification for each of the curtained regions.  
23 Alternatively, a single specification can explicitly designate the ring or subring  
24 that each address range in the specification belongs to.

1        If the current instruction address from pointer 308 matches any of the  
2        addresses in map 310, that instruction is included in a particular curtained code  
3        ring or module. Curtain logic 314 then permits the control unit to issue signals  
4        298 for performing certain operations, including reading and writing memory  
5        locations in the same ring, or a less privileged ring that might contain secrets.  
6        (Additionally, as described below, certain opcodes are restricted to executing only  
7        when the CPU is executing curtained code.) For example, if decoder 304 is  
8        executing an instruction not located within the range of curtained memory, and if  
9        that instruction includes an operand address located within the curtained-memory  
10       specification, control unit 296 blocks the signals 298 for reading the data at that  
11       address and for writing anything to that address. If a non-privileged access is  
12       attempted, the CPU or memory system can flag an error, fail silently, or take other  
13       appropriate action. If it is desired to place the curtain logic on a chip other than  
14       the processor, a new microprocessor instruction or operating mode can strobe the  
15       instruction pointer's contents or other information related to the code executing  
16       onto an external bus for comparison with the curtained address ranges.

17       The execution of trusted code routines is frequently initiated by other  
18       programs that are less trusted. Therefore, curtain logic 314 must provide for some  
19       form of execution access to the curtained code stored in Rings A and B. However,  
20       full call or jump accesses from arbitrary outside code, or into arbitrary locations of  
21       the curtained memory regions, might possibly manipulate the secure code, or  
22       pieces of it, in a way that would reveal secret data or algorithms in the curtained  
23       memory. For this reason, logic 314 restricts execution entry points into curtained  
24       memory regions 256 and 258 as well as restricting read/write access to those  
25       regions. In one embodiment, the curtained code exposes certain entry points that

1 the code writers have identified as being safe. These often occur along functional  
2 lines. For instance, each operation that a piece of curtained code can perform has  
3 an accompanying entry point. Calling subroutines at these entry points is  
4 permitted, but attempts to jump or call code at other entry points causes an  
5 execution fault.

6 An alternative allows automated checking of entry points and provides  
7 additional granularity of rights by permitting entry to curtained memory functions  
8 only through a special entry instruction. For example, a new curtained-call  
9 instruction, CCALL Ring, Subring, OpIndex, has operands that specify a  
10 ring, a subring, and a designation of an operation whose code is located within that  
11 ring and subring. This instruction performs conventional subroutine-call  
12 operations such as pushing a return address on a stack and saving state  
13 information. The stack or the caller's memory can be used to pass any required  
14 parameters. A conventional RETURN instruction within the curtained code returns  
15 control to the calling routine. Return values can be placed in memory, registers,  
16 etc.

17 When decoder 304 receives a CCALL instruction, curtain entry logic 314  
18 determines whether the calling code has the proper privileges, and whether the  
19 instruction's parameters are valid. If both of these conditions are met, then the  
20 instruction is executed and the curtained routine is executed from its memory ring.  
21 If either condition does not hold, logic 314 fails the operation without executing  
22 the called code.

23 Logic 314 determines whether or not to execute the code by comparing the  
24 privilege level of the calling code and the operation-index parameter, with entries  
25 in a jump-target table 316 stored in a location accessible to it. The logic to enforce

1 these requirements can be implemented in the memory controller 314, or by code  
2 executing in a highly privileged ring such as Ring A. Table I below illustrates one  
3 form of jump-target table. The table can be stored in the same curtained memory  
4 block as the code itself, or in a memory block that is more privileged; or it can be  
5 stored in special-purpose storage internal to the CPU or memory manager.

6  
7 Table I  
8

Index	Target Address	User	Kernel	Curtain
0	BAB-PC	FALSE	TRUE	TRUE
1	REVEAL-PC	TRUE	TRUE	TRUE
2	LOAD-PC	FALSE	FALSE	TRUE

10  
11 An entry for each index, 0-2, gives the (symbolic) target or start address of  
12 the code for that operation, and the privileges levels—user, kernel, or curtained—  
13 that are permitted to execute the code. “Curtained” level means that only other  
14 curtained code can call the routine. Other or finer privilege levels are possible.  
15 As an alternative to the above jump table, entry logic 314 could permit only a  
16 single entry point into each ring of curtained memory, and employ a passed  
17 parameter to specify a particular operation. Or it could, for example, permit calls  
18 only to addresses that are predefined as the beginnings of operations. The  
19 curtained code itself could verify and call the operation.

20 Restricting call access to curtained code within processor 282 still leaves  
21 open the possibility that outside rogue programs or devices might be able to hijack  
22 the code after its execution has begun in order to obtain secrets left in registers, or  
23 to otherwise modify machine state to subvert operation. Therefore, control unit  
24 296 ensures atomicity in executing the curtained code: once started, the code  
25 performs its entire operation without interruption from any point outside the

1 secure curtained-memory regions. In many cases, it is not necessary to execute an  
2 entire function atomically, but only a part. For example, only the code that  
3 verifies a bus-master card's identity need be performed atomically, and not its  
4 total initialization module. Additionally, the entry code could take steps to prevent  
5 subversion. For instance, it could establish new interrupt vectors and memory  
6 page tables. In this case, only the entry code must be executed atomically. Once a  
7 new interrupt table is established, trusted code can ensure that secrets do not leak  
8 in response to an interrupt.

9 Fig. 9 is a flowchart 332 illustrating an exemplary process for providing  
10 curtained execution protection in a processor such as processor 282 of Fig. 8. The  
11 process of Fig. 9 is implemented by a processor, such as processor 282. Fig. 9 is  
12 described with additional reference to components in Figs. 7 and 8

13 For a single-level curtained memory, process 332 refers to the entire  
14 curtained region. For a memory organization such as 252 of Fig. 7 having  
15 multiple rings or subrings, the term "curtained region" means the levels inside or  
16 beside the ring in which the current instruction is located. For example, the  
17 curtained region for an instruction whose address is in Ring C in Fig. 7 comprises  
18 Ring B (including all its subrings) and Ring A; the curtained region for an  
19 instruction in Subring B1 comprises Subrings B2 and B3 (but not the rest of Ring  
20 B) and Ring A.

21 After the current instruction is decoded (step 334), memory addresses  
22 associated with the instruction are tested. If the instruction uses virtual addresses,  
23 the tests operate upon the physical addresses as translated in step 334. A  
24 determination is made whether the instruction accesses any memory location  
25 during its execution (step 336). An instruction might read an operand or write

1 data to a memory address, for example. If the instruction does not access any  
2 memory, or at least any memory that might contain a curtained region, then the  
3 instruction is executed (step 338). If the instruction does involve a memory  
4 location, the address is tested to determine whether it is within a region that is  
5 curtained off from the current region (step 340). If not, the instruction is executed  
6 (step 338). If so, the type of access requested by the instruction is determined  
7 (step 342). If the access is anything other than the special curtained-call opcode,  
8 then a fault is signaled (step 344), and an appropriate error routine or logic circuit  
9 blocks the access. Other accesses include reading data from the location, writing  
10 data to it, or executing a normal instruction there.

11 The only access permitted into a curtained-memory ring is an execution  
12 access by a particular kind of instruction, such as the curtained call (CCall)  
13 discussed above. If an instruction is detected in step 342 as desiring to initiate  
14 execution of code at a location inside a region curtained from the current region, a  
15 determination is made as to whether the target entry point is valid (step 346) —  
16 that is, whether the requested index is in the jump table. A determination is also  
17 made as to whether the current instruction has the privilege level required to  
18 invoke the operation at the desired location (step 348). If either test fails (step 346  
19 or 348), a fault is signaled (step 344). If both pass, the curtain-call instruction as  
20 described above is executed (step 350).

21 Steps 352-356 navigate among the rings and subrings of the curtained  
22 memory. A CCALL instruction causes the curtained-memory ring containing the  
23 target address of the call to be opened (step 352). That is, it makes that ring the  
24 current ring for the purposes of process 332. A routine starting at that address thus  
25 has read/write and execution access to the memory of the ring, and only rings

1 inside or peer to that ring are now restricted curtained memory. Also in step 352,  
2 any extra protection for ensuring atomicity of the routine being executed at the  
3 new current level, such as interrupt suspension or bus locking, is engaged. A  
4 routine executing in curtained memory can end with a normal Return  
5 instruction. If the routine was called from a less secure ring, step 354 causes the  
6 current ring to close (step 356) and retreat to the ring from which the call was  
7 made, either a less secure ring of curtained memory, or the outer, unsecured  
8 memory of Ring C.

9 Fig. 10 details an example of hardware 370, added to computer 118 of Fig.  
10 2, for defining curtained code as a set of secure memory pages or regions within  
11 the same address space as that of the system memory. Hardware 370 can be  
12 incorporated into a system controller chip of the chipset, or directly into a  
13 processor.

14 An 8-bit curtained-code (CC) register 372 holds a designation of one of a  
15 number of security managers running in Ring B, and a designation of one of a  
16 number of applications running in Ring B. In general, CCR=n,m signifies  
17 application m running in Ring B under the control of security manager n in Ring  
18 B. CCR=n,0 means that security manager n itself is executing. CCR=0,0  
19 indicates that a program is running in Ring A, and CCR=FF indicates code  
20 running in the unprotected Ring C. The contents of the CCR thus indicate which  
21 program is currently running and in what ring or security level. A system having  
22 multiple processors has a separate CCR for each one.

23 Access-control table (ACT) 374 is a fast read/write memory containing  
24 64K entries each associated with a page of memory in system memory 141,  
25 specified by bits A16:23 of address bus segment 376, and with a certain

1 combination of Ring-B programs specified by CCR 372. Each ACT entry  
2 contains bits that determine the rights for a program combination that accesses that  
3 page. Typically, one bit of each entry indicates whether the programs specified by  
4 the CCR contents has read privileges for the page specified. Another bit indicates  
5 write privileges for the page. A third bit can indicate execution privileges, that is,  
6 whether the programs are permitted to jump to or call a program within the page.

7 The ACT entry contents cause gate 378 to pass or to block processor read  
8 and write control lines 380 and 382 to memory read and write lines 384 and 386.  
9 In addition, a signal on line 388 indicates that the processor is about to execute  
10 program code within the page. (Although most present-day microprocessors do  
11 not emit such a signal directly, it can be generated easily.) The third ACT bit  
12 controls memory read line 384 to allow or to block reading the addressed byte into  
13 the processor for execution. The address layout shown in Fig. 10 envisions a  
14 secure memory region having a total of 256 pages addressed by bus segment 376.  
15 Each page has 64K bytes, addressed by bits A00:15 of low-order bus segment 388.  
16 High-order address bus segment 390 places these pages at a certain point within  
17 the overall address space of system memory by enabling block 392 only for a  
18 certain combination of address bits A24:31. Block 392 allows gate 378 to pass  
19 signals 380 and 382 directly to memory lines 384 and 386 except for addresses  
20 within the secure memory region.

21 Connections from system data bus 394 permit the contents of register 372  
22 and access-control table 374 to be modified in a secure mode when a particular  
23 port number is addressed by the processor. Programming these components will  
24 be described later.

1 Computer 118 calls code in a secure page using a special trap facility that  
2 can be located in a processor or bus-control chipset, or in memory-manager logic  
3 for legacy systems. In this implementation, only a single entry point serves all  
4 code executing in any secure page. A calling program identifies the particular  
5 routine to be executed as a parameter in a special call instruction.

6 In addition, a “security porch” enhances execution security for older  
7 processors. A memory controller watches instruction fetches after a secure-page  
8 routine is called to ensure that the next sixteen or so fetches are executed in  
9 ascending order by the initiating processor. If this is the case, then this porch code  
10 can be designed so that its keys and self-security instructions are mapped into  
11 memory in the region following the porch code for the initiating processor. For  
12 example, porch code can set the interrupt vector to a secure handler in secure  
13 memory and switch off caching for certain regions, or invalidate the cache. Once  
14 this is accomplished, the curtained code can execute without fear of preemption by  
15 untrusted code. Enforcing an entry point and protecting the first few dozen  
16 instructions allows the curtained memory to be visible to other calls without  
17 compromise. An alternative implementation could use a trap into memory  
18 protected by a processor system mode, in which the processor is protected against  
19 interruption; the system-mode handler can then turn off interrupts and jump into  
20 the protected code rather than returning to the interrupting program. However,  
21 further protection would still be required against multi-processor or other bus-  
22 master attacks.

23 Fig. 11 diagrams components 400 employed in the secure pages approach  
24 in an illustrative embodiment. Adaptations for other forms of data are within the  
25 art. In Fig. 11, software modules running in and storing data in secure memory

1 pages are shown as blocks with rounded corners. Arrows labeled "trust"  
2 symbolize the granting of trust from a module to one or more other modules.  
3 Dashed horizontal lines indicate processor operating modes.

4 The root of trust of secure code modules is a secure loader 402, running in  
5 Ring A of the secure memory region, under a secure system mode of the  
6 processor. Trust descends hierarchically to one or more security managers 404  
7 executing in Ring B, or in one of its subrings, in kernel mode, the same mode as  
8 the rest of OS 123. As shown in Fig. 11, most of the remainder of OS 123 need  
9 not have any additional protection. Security manager 404 in turn confers trust  
10 upon a memory manager responsible for a group of pages of the secure memory;  
11 manager 404 also runs in Ring B. Security manager 404 also confers trust at a  
12 lower level upon certain designated third parties. A banking institution can  
13 provide an application program such as 406 for allowing a user of the computer to  
14 access his or her bank account(s) or purchase goods and services. In this example,  
15 only one dynamic link library (DLL) 408 used by application 406 requires trust in  
16 order to protect the user's private data (e.g., account numbers, balances,  
17 transactions, etc.). In fact, because DLLs are frequently shared among multiple  
18 programs, it is possible that a single secure DLL can provide protection for other  
19 applications as well. Module 408 runs in the processor's user mode, in Ring B of  
20 the secure memory. The remaining applications, as well as all other non-secure  
21 modules in Fig. 11, run in the outermost memory Ring C. Alternatively, all of the  
22 application 406 may require trust and be running in Ring B of the secure memory.

23 The secure module 408, receiving trust from security manager 404, can in  
24 turn entrust a further, lower level of trust in other modules. Each secure module  
25 can contain names of one or more other specific modules it is willing to trust, and

1 can confer trust upon them when it receives trust from higher in the tree. Trust  
2 can be established declaratively or programmatically, by naming the public key,  
3 digest, or other signature of other applications or modules that have access to some  
4 or all of their code or data in the secure storage. Security manager 404, via its  
5 privileged interaction with memory manager 412, interprets and establishes these  
6 hierarchical trust relationships.

7 Security manager 404 also provides a way for the system to vouch  
8 cryptographically for the digest or signature of code running in a secure page.  
9 When computer 118 is communicating with a portable IC device, such as portable  
10 IC device 116 of Fig. 1, SM 404 signs a certificate that a particular application  
11 component (such as DLL 408 or application 406) is running in a secure page. The  
12 portable IC device thus knows whether the application or module that will handle  
13 the sensitive data is operating in a secure manner. As long as neither the SM 404  
14 nor the hardware security has been compromised, the portable IC device knows all  
15 the information necessary to establish trust, namely, that other components in the  
16 operating system cannot access the private data. The portable IC device may also  
17 name other components that it trusts to work cooperatively on the private data.  
18 The portable IC device therefore need only trust certain applications to receive the  
19 private data, and the SM 404 and security hardware will ensure that the application  
20 is not subverted and that its secrets are not read or changed by untrusted code.

21 Code running in the secure memory pages cooperates with the code in  
22 other, untrusted modules of the main operating system 123. As described above,  
23 any module can call into a secure page using the special trap facility. The secure-  
24 page code can exit normally back to an untrusted module, relinquishing its  
25 protected status as it does so; or, a timer event or other interrupt can occur upon

1 exit. In order to reduce changes to the remainder of the OS, its normal interrupt  
2 handler 401 is permitted to field interrupts, even though it is not trusted.  
3 Therefore, security manager 404 sets up a new trusted interrupt handler 410. This  
4 handler saves the system state in a secure page, so that no other code can access it  
5 or any secrets it might contain. Handler 410 can then initiate a software interrupt  
6 that the untrusted part of the OS is allowed to process. Normal OS mechanisms  
7 reschedule security manager 404, which in turn reschedules the secure module.

8 The security manager 404 is responsible for setting up an address space in  
9 protected memory appropriate for a secure application or module that is about to  
10 run. For instance, when the OS context switches into a trusted kernel component,  
11 a banking application's own secure pages, all the rest of kernel memory, and any  
12 other modules that have explicitly granted trust on this OS module will be mapped  
13 into active memory. It should be noted here that the mode of the processor is tied  
14 directly to the code that is executing. For instance, there is no 'enter user mode'  
15 instruction; the processor merely traps to the code entry points in a secure page.  
16 During this transition, the mode changes automatically. This allows placing secret  
17 keys and code in a module from which only the trusted code can ever access the  
18 secret keys, providing a broad general-purpose model of platform security.

19 Security manager 404 is the only piece of code that is allowed to set or  
20 change page permissions in secure memory. The SM 404 establishes a memory  
21 map by telling the programmable memory manager to map or hide various pages  
22 in the system-memory address space by writing to a port or a memory address that  
23 programs memory manager 412.

24 When an OS module calls the SM 404 indicating an application 406 that  
25 should execute, the SM 404 maps the application's pages, along with any other

1 secure pages that the application has access to, into the processor's address space,  
2 then starts executing the application. Should an interrupt occur, the trusted code  
3 saves the current machine state in a safe place. The security manager allows OS  
4 111 to handle the interrupt. The normal interrupt handler 413 reschedules the SM  
5 404, which later restarts the interrupted application with the safely saved state.  
6 Additionally, code in a secure page might have to call functions provided by an  
7 untrusted module of the operating system. It can manage this by requesting  
8 security manager 404 to save its state in a safe location before calling the normal  
9 OS module, then restoring the state and marshalling parameters back into the  
10 secure page when the OS module returns. Code in a secure page has access to all  
11 other virtual memory. For instance, a kernel secure-page component will have  
12 access to all of kernel memory in addition to its own memory and some other  
13 selected secure pages. Further, a user-mode secure-page component has the  
14 application's memory mapped into its virtual memory, and runs in user mode.

15 Security-manager handler 414 runs as a normal OS module with kernel  
16 privileges but no special security or trust. It has two broad sets of functions. The  
17 first is to provide an interface to the OS for loading and unloading curtained code  
18 for secure modules. The second is to provide the normal-code entry point to SM  
19 functions. It is called on an OS thread and handles the context switch into the  
20 secure code comprising the SM. Should the SM be preempted, the SM-Handler  
21 will trampoline the interrupt to the OS interrupt handler 401, and will in turn be re-  
22 scheduled by the OS after the interrupt has been handled.

23 Another function of security manager 404 is to offer several types of  
24 cryptographic services to the application modules. First, as mentioned, it offers a  
25 way for an application to name other applications that it trusts to access its data.

1 As a practical matter, this can be accomplished by placing binary resources in an  
2 application-level secure DLL such as 408 that name the digest or public key of  
3 other trusted modules. Upon initial load—or later—the SM 404 identifies the  
4 trusted modules, so that a call to them maps their pages. The second facility  
5 allows a user to establish cryptographically which application he is  
6 communicating with. There are several ways of doing this. A convenient way is  
7 to allow the user's portable IC device to encrypt a data block that contains its keys  
8 or other secret data, and that names the digest of the target secure application. The  
9 application alone is able to decrypt the secrets, and only gives the secrets to the  
10 named application. The third facility is a local facility in computer 118 for storing  
11 secret data such that only it or another trusted component can access them. A  
12 convenient way of doing this is to allow the application to store a secret of a target  
13 application encrypted with a key provided by SM 404. The SM 404 then ensures  
14 that only the named application ever gets to access the secret.

15 Because a security manager is relatively complex and is specific to a  
16 particular operating system, it is typically ill suited to being hard-coded in the  
17 processor or motherboard of a system by the manufacturer. On the other hand,  
18 providing sufficient security usually requires some form of protection at a system  
19 level. This embodiment employs a small, generic code module fixed in ROM at  
20 the manufacturer level to provide a small number of generic functions adaptable  
21 for use with many different security managers. Secure loader 402 loads an OS-  
22 specific security manager 404 into a secure page and vouches for it  
23 cryptographically, using the loader's own keys or secrets in a secure online  
24 session. Initially, the Secure Loader might even receive some code and secrets  
25 online from a system manufacturer, OS vendor, or internet service vendor. After

1 this one online interaction, the security manager can be stored in encrypted form  
2 on disk, and its decryption and verification managed by the secure loader when the  
3 system is booted. If the Security Manager does not contain secrets—e.g., if it  
4 relies upon certificates minted by the secure loader—it could even be shipped in  
5 cleartext on a CD-ROM.

6 Secure loader 402 typically conducts an authenticated encrypted online  
7 session with an OS vendor to establish that the vendor can ensure that the code  
8 and keys are being issued to a trusted module. Encryption is needed to ensure that  
9 untrusted parties, including the owners and legitimate users of the system, cannot  
10 modify the code or steal embedded secrets. The OS vendor transmits appropriate  
11 SM code and unique keys and a certificate that can be used in subsequent boots.  
12 The security loader stores the SM to designated secure pages, and also stores it to  
13 disk or other permanent storage in encrypted form so that it can be used  
14 subsequently without an online step.

15 When security manager 404 receives a request to execute a particular  
16 application such as 406, it restricts the pages of secure memory that can be  
17 mapped. In this implementation, security manager 404 informs memory manager  
18 412 which mode it is about to enter. The memory manger then adjusts access-  
19 control table 374 of Fig. 10 to contain the proper per-page access permissions, i.e.,  
20 the values of the read, write, and execute bits for each page of secure memory.  
21 Again, the memory manager can only restrict the permissions initially  
22 programmed into the table 374 by security manager 404, and cannot expand them.  
23 When an application module completes execution, it maps its pages out of  
24 memory, or requests the security manager to do so. It then relinquishes control.

1 As mentioned earlier, memory manager 412 must be programmed with the  
2 various privileges for each page of secure memory at the different levels specified  
3 by different possible contents of register 372 of Fig. 10. It is convenient here to  
4 make programming the access-control table possible only when secure loader 402  
5 is executing, that is, when CCR=0,0 indicates that execution is taking place in  
6 Ring A. Programming can be accomplished by writing to a dedicated port or  
7 mapping it into memory and issuing ordinary reads and writes. The secure loader  
8 has its own rudimentary software memory manager that allocates unused pages to  
9 the Ring B security managers 404 on demand. These specify the protections to be  
10 applied to the pages they own, and the Ring A code programs the hardware table  
11 374 accordingly. Of course, loader 402 must not be subject to spoofing attacks;  
12 therefore, the security managers 404 pass parameters in their own secure memory,  
13 so that another attacking SM, or the OS, cannot change the request.

14 Fig. 12 is a flowchart illustrating an exemplary process for activating the  
15 secure loader. Description of the steps of method 420 in a particular order does  
16 not necessarily imply any specific temporal sequence. The process of Fig. 12 is  
17 implemented by a computer (e.g., computer 118 of Fig. 2), and may be performed  
18 partially or wholly in software. Fig. 12 is described with additional reference to  
19 components in Figs. 10 and 11.

20 When the computer starts, secure loader 402 of Fig. 11 is activated (steps  
21 422 – 434). The processor is reset in a high-security system mode (step 422), and  
22 the loader is booted into Ring A of secure memory in this mode (step 424). An  
23 encrypted connection to the OS vendor is initiated (step 426), and an  
24 authentication (such as a signed certificate) is minted and issued to the vendor  
25 (step 428). If the vendor decides not to trust the system, then the security manager

1 module is not loaded and the process ends. However, if the vendor decides to trust  
2 the system (step 430), it downloads an encrypted security manager module 404  
3 and keys (step 432), and stores it to disk (step 434). (The vendor or others can  
4 download other code as well, if desired.) Encryption assures that untrusted  
5 persons, including the owner of computer 118, cannot modify the code or steal  
6 secret keys or other downloaded data. If it does not itself contain any secret data,  
7 and relies on verifications supplied by loader 402, then manager 404 can be stored  
8 in a non-encrypted cleartext form. Storing the code and data in a nonvolatile  
9 storage obviates the need for downloading the same code and keys again when  
10 system 118 is rebooted. Alternatively, the security manager can be supplied with  
11 the OS on a CD-ROM or other medium so that it does not need to be downloaded.  
12 Fig. 12 shows only a single security manager. Different OSs, or even the same  
13 OS, can have multiple security managers.

14 In 436 - 440, secure loader 402 programs memory manager 412. Each page  
15 of memory (step 436) has a set of permissions for different CC-register contents  
16 (step 438). Again, CC register 372 defines which secure-memory ring and subring  
17 is active. The read, write, and execute permissions for each entry of access-  
18 control table 374 are defined (step 440). Modules farther down in the trust  
19 hierarchy can restrict these privileges, but cannot expand them.

20  
21 **Operation**

22 Fig. 13 is a flowchart illustrating an exemplary process for two-way  
23 authentication in accordance with the invention. The process of Fig. 13 is  
24 implemented by a combination of a portable IC device (e.g., device 116 of Fig. 2)  
25 and a computer (e.g., computer 118 of Fig. 2), and may be performed partially or

1 wholly in software. Fig. 13 is described with additional reference to components  
2 in Fig. 2.

3 Communication between the portable IC device 116 and the computer 118  
4 is initially established (step 462). Applications executing on the computer can be  
5 authenticated, such as via the authenticated boot or curtaining methodologies  
6 discussed above. The communication can be established by communicatively  
7 coupling the portable IC device to the computer using any of a variety of  
8 conventional coupling technologies.

9 Once such communication is established, a secure channel between the  
10 portable IC device and the application is established (step 464). The secure  
11 channel provides an assurance that others cannot intercept, modify, replay, or  
12 decipher messages being exchanged between the IC device and the application via  
13 the channel. A key-exchange protocol such as the well-known Diffie-Hellman  
14 key-agreement protocol is used to establish the secure channel in step 464.  
15 Alternatively, other conventional cryptographic techniques can be used to establish  
16 the secure channel between the portable IC device and the computer.

17 Once a secure channel between the portable IC device and the application is  
18 established, the application and optionally the OS executing on the computer  
19 authenticate themselves to the portable IC device (step 466). This application is  
20 the application that will be accessing information on the portable IC device. For  
21 example, the application may be a banking application that will be ordering  
22 products or services as requested by the user and using a signature provided by the  
23 portable IC device.

24 Whether trusted communication between the portable IC device and the  
25 desired application and OS can proceed is dependent, in part, on whether the

1 portable IC device can verify the authenticity of the application (step 468). If the  
2 portable IC device cannot verify the authenticity of the application, OS, and  
3 hardware, then trusted communication cannot proceed (step 470). It should be  
4 noted that if the portable IC device cannot verify the authenticity of the  
5 application, etc., trusted communication cannot proceed regardless of whether the  
6 portable IC device is authenticated to the application. If the desired application  
7 and OS can indeed authenticate themselves to the portable IC device, the portable  
8 IC device indicates this fact to the user via the optional indicator light 132 or  
9 similar mechanism. The user then knows that it is safe to trust the computer.

10 If the portable IC device can verify the authenticity of the application, then  
11 the portable IC device authenticates itself to the application (step 472). The user  
12 can use the facilities of the computer to do so, such as the keyboard and display,  
13 since they are under the control of a trusted application and OS. Whether trusted  
14 communication between the portable IC device and the application can proceed is  
15 also dependent on whether the application can verify the authenticity of the  
16 portable IC device (step 474). If the application cannot verify the authenticity of  
17 the portable IC device, then trusted communication cannot proceed (step 470).  
18 However, if the application can verify the authenticity of the portable IC device,  
19 then trusted communication can proceed (step 476), with both the portable IC  
20 device and the computer knowing that the other is trustworthy. An application that  
21 is deemed trustworthy provides an assurance to the user that the application will  
22 not perform unauthorized actions on behalf of the user (e.g., order products using  
23 the user's signature). The portable IC device can thus unlock itself and make the  
24 private information stored thereon accessible to the application.

25

1 In the illustrated example, the computer authenticates itself to the portable  
2 IC device (step 466) before the portable IC device authenticates itself to the  
3 computer (step 472). By having the computer authenticate itself first, the user can  
4 trust the computer to accept a personal identification number (PIN) or other  
5 confidential information used by the portable IC device to authenticate itself to the  
6 computer. Alternatively, these authentication steps 466 and 472 could occur in  
7 reversed order (with the portable IC device authenticating itself to the computer  
8 first) or substantially concurrently.

9 Fig. 14 is a flowchart illustrating an exemplary process for a computer  
10 application to authenticate itself to a portable IC device in accordance with the  
11 invention. The process of Fig. 14 is implemented by a combination of a portable  
12 IC device (e.g., device 116 of Fig. 1) and a public computer (e.g., computer 102 or  
13 104 of Fig. 1), and may be performed partially or wholly in software. Fig. 14 is  
14 described with additional reference to components in Figs. 2, 3, and 11, and  
15 describes step 466 of Fig. 13 in more detail.

16 The application that will be accessing information on the portable IC device  
17 initially sends a request to the portable IC device to unlock itself (step 502). The  
18 portable IC device responds to the request by sending a challenge nonce to the  
19 application (step 504). In the illustrated example, the challenge nonce comprises a  
20 random number generated by the portable IC device. Upon receiving the  
21 challenge nonce, the application responds to the portable IC device (step 506). In  
22 the illustrated example, the response is generated by signing the received random  
23 number using a private key for the application.

24 The portable IC device then verifies the response (step 508). The response  
25 is verified using the public key for the application. If the response cannot be

1 verified, then the application is not verified and trusted communication between  
2 the portable IC device and the application cannot occur. However, if the response  
3 can be verified, then the portable IC device sends a definition of a set of trusted  
4 applications to the application (step 510). This definition of a set of trusted  
5 applications can be an explicit, pre-determined list that is maintained on the  
6 portable IC device, or it can be a list of rules maintained on the portable IC device  
7 that implicitly define such a list of trusted applications. This list can be provided  
8 by the manufacturer of device 116, or alternatively can be generated (or modified)  
9 by an administrator, the distributor of the device 116, or the user. Examples of  
10 such applications include banking applications, applications controlling  
11 purchasing for retailers, etc. Alternatively, the list may only identify trustworthy  
12 processor(s) and operating system(s) with the individual applications relying on  
13 the operating system to vouch for their trustworthiness.

14 Upon receiving the definition of the set of trusted applications, the  
15 application has a certificate chain generated indicating that it is one of the trusted  
16 applications and is currently executing on the computer (step 512). The certificate  
17 chain authenticates the application – proving that the trusted application was  
18 running on the computer when the challenge was received. This certificate chain  
19 is then sent to the portable IC device (step 514). The portable IC device then  
20 verifies the certificate chain received from the application (step 516). Trusted  
21 communication between the application and the portable IC device can only  
22 proceed if the certificate chain can be verified.

23 The manner in which the certificate chain is generated can vary, depending  
24 on the authentication methodology being used. To generate the certificate chain  
25 using the authenticated boot methodology, CPU 134 of Fig. 3 mints an OS

1 certificate that contains data from the portable IC device 116 and an identity of the  
2 OS. The OS certificate takes the following form:

3  
4 OS Certificate = (SIR, Reply, Data,  $K_{CPU}^{-1}$ ) signed by  $K_{CPU}^{-1}$   
5

6 The “data” can be the challenge nonce from step 504. In addition to the  
7 data, the OS certificate contains the SIR value, a reply, and the CPU’s public key  
8  $K_{CPU}$ . The “reply” can optionally contain all of the data written to the boot log 70  
9 so that the portable IC device can evaluate what software components are  
10 currently loaded and executing. In other cases, the portable IC device could just  
11 trust the OS publisher (and hence simply the value of the SIR). The OS certificate  
12 is signed using the CPU’s private key  $K_{CPU}^{-1}$ . Effectively, the OS certificate says  
13 “The processor named  $K_{CPU}$  was running the Operating System SIR with the  
14 specified boot log when it received the challenge”. (In the case where the boot log  
15 is included, the CPU is including more information, effectively saying “Further, it  
16 was running this OS revision, with these version components, and these device  
17 drivers and applications.”)

18 The newly-minted OS certificate and the CPU manufacturer’s certificate  
19 152 and OEM certificate 141 are returned to the portable IC device as the  
20 certificate chain. The OS certificate and manufacturers’ certificates are validated  
21 using a series of tests. Failure of any one of the tests results in failure of the  
22 verification step 516. However, passing all tests authenticates the application to  
23 the portable IC device, proving that the trusted application is running on the  
24 computer 118.

1        The first test is whether the portable IC device recognizes the SIR value  
2        contained in the OS certificate and trusts the associated operating system.

3        Assuming the OS is trusted, the portable IC device next determines whether  
4        the data is the same challenge nonce that it generated and supplied to the  
5        application. If the data returned in the reply fails to match the challenge provided  
6        by the portable IC device, the verification fails. However, if the two match, the  
7        portable IC device evaluates whether the OS certificate is properly signed with the  
8        CPU's private key  $K_{CPU}^{-1}$ . The portable IC device makes this evaluation using the  
9        enclosed public key  $K_{CPU}$ .

10       With respect to the CPU manufacturer's certificate, the portable IC device  
11       determines whether the certificate names the same public key  $K_{CPU}$  used in the OS  
12       certificate. If so, the portable IC device continues to the next test; otherwise, the  
13       verification fails.

14       The portable IC device next examines whether the manufacturer certificate  
15       is signed by the manufacturer's private key  $K_{MFR}^{-1}$  by using the manufacturer's  
16       public key  $K_{MFR}$  and whether the OEM certificate is signed by the OEM's private  
17       key  $K_{OEM}^{-1}$  by using the OEM's public key  $K_{OEM}$ . If the signatures are not proper,  
18       the verification fails.

19       If the signatures are proper, the portable IC device verifies that the  
20       application it is communicating with is indeed a trusted application. If the  
21       application is not on the list of trusted applications, then the verification fails.

22       Alternatively, rather than having the portable IC device evaluate whether a  
23       trusted application is connected, the evaluation could be performed by the trusted  
24       operating system. The portable IC device would provide the definition of the list  
25       of trusted applications to the application, which would make the definition

1 available to the operating system. The operating system certificate and the  
2 processor and OEM certificates would be provided analogous to the above  
3 discussion, however, the “reply” could include an indication of the application that  
4 is connected to the portable IC device (the same application that requested the  
5 portable IC device to unlock itself). Additionally, the portable IC device can use  
6 the boot log to determine whether any untrusted OS components have been  
7 loaded.

8 In order to generate the certificate chain using the curtaining methodology,  
9 the process is similar to that of the authenticated boot methodology, however a  
10 certificate from the operating system is not necessary. Rather, the security  
11 manager 404 of Fig. 11 would mint a new certificate taking the following form:

12  
13 Certificate = (Data,  $K_{\text{manager}}$ ) signed by  $K_{\text{manager}}^{-1}$

14  
15 The “data” includes the challenge nonce from step 504.

16 (In this formulation, the security manager holds  $K_{\text{manager}}^{-1}$  in sealed storage.  
17 This approach eliminates the need for transmitting the processor certificate and  
18 OEM certificate, since they must already have been used to remove  $K_{\text{manager}}^{-1}$  from  
19 sealed storage.)

20 The newly minted loader certificate is returned to the portable IC device as  
21 the certificate chain. The certificates are validated by the portable IC device using  
22 a series of tests, failure of any one of which results in failure of the verification  
23 step 516. The first test is whether the data identifies a trusted application as being  
24 connected to the portable IC device. The second test is whether the certificate is  
25

1 properly signed with the security manager's private key  $K_{\text{manager}}^{-1}$ . The portable IC  
2 device makes this evaluation using the known public key  $K_{\text{manager}}$ .

3 Fig. 15 is a flowchart illustrating an exemplary process for a portable IC  
4 device to authenticate itself to a computer system in accordance with the  
5 invention. The process of Fig. 15 is implemented by a combination of a portable  
6 IC device (e.g., device 116 of Fig. 1) and a public computer (e.g., computer 102 or  
7 104 of Fig. 1), and may be performed partially or wholly in software. Fig. 15 is  
8 described with additional reference to components in Fig. 2, and describes step  
9 472 of Fig. 13 in more detail.

10 The application that will be accessing information on the portable IC device  
11 initially sends a challenge, also referred to as a "challenge nonce", to the portable  
12 IC device (step 522). In the illustrated example, the challenge nonce comprises a  
13 random number generated by the application. Upon receiving the challenge  
14 nonce, the portable IC device responds to the challenge (step 524). In the  
15 illustrated example, the response is generated by signing the received random  
16 number using the portable IC device's private key. This signed number is then  
17 returned to the application as the response.

18 Upon receiving the response, the application verifies the response (step  
19 526). In the illustrated example, the response is verified using the portable IC  
20 device's public key, which is known to the application. The public key can be  
21 made known to the application in any of a variety of conventional manners, such  
22 as transmitting the public key to the application when communication between the  
23 application and the portable IC device is initially established (step 462 of Fig. 13).  
24 As only the portable IC device knows the portable IC device's private key, the  
25 application can verify the authenticity of the portable IC device by evaluating,

1 using the portable IC device's public key, whether the random number was  
2 properly signed with the portable IC device's private key.

3 Additional user-verification may also be required, such as requiring the user  
4 to enter a valid PIN. This user-verification may be implemented as part of the  
5 response and verification steps 524 and 526, or alternatively may be an additional  
6 set of steps after the response is verified in step 526.

7

8 **Conclusion**

9 Thus, the invention provides for authenticating an open system application  
10 to a portable IC device. This authentication allows the authenticity of an  
11 application(s) on the open system to be proven to the portable IC device. The  
12 authentication advantageously allows the application(s) on the open system to be  
13 trusted by the portable IC device, providing an assurance that the private nature of  
14 information on the portable IC device made available to the application(s) will be  
15 maintained and that such information will not be misused.

16 Although the invention has been described in language specific to structural  
17 features and/or methodological steps, it is to be understood that the invention  
18 defined in the appended claims is not necessarily limited to the specific features or  
19 steps described. Rather, the specific features and steps are disclosed as preferred  
20 forms of implementing the claimed invention.